UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/826,888	04/16/2004	Sitaram Dontu	CIS0216US	1421

33031 7590 11/10/2009

CAMPBELL STEPHENSON LLP
11401 CENTURY OAKS TERRACE
BLDG. H, SUITE 250
AUSTIN, TX 78758

EXAMINER	
REDDIVALAM, SRINIVASA R	
ART UNIT	**PAPER NUMBER**
2477	
MAIL DATE	**DELIVERY MODE**
11/10/2009	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	10/826,888	DONTU ET AL.	
	Examiner	Art Unit	
	SRINIVASA R. REDDIVALAM	2477	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *17 June 2009*.

2a)☒ This action is **FINAL**. 2b)☐ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1,3-34,36-43,45-57 and 59-70* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1,3-34,36-43,45-57 and 59-70* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☐ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☒ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date *06/18/2009*.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____.
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

Claim Rejections - 35 USC § 103

1. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

2. This application currently names joint inventors. In considering patentability of the claims under 35 U.S.C. 103(a), the examiner presumes that the subject matter of the various claims was commonly owned at the time any inventions covered therein were made absent any evidence to the contrary. Applicant is advised of the obligation under 37 CFR 1.56 to point out the inventor and invention dates of each claim that was not commonly owned at the time a later invention was made in order for the examiner to consider the applicability of 35 U.S.C. 103(c) and potential 35 U.S.C. 102(e), (f) or (g) prior art under 35 U.S.C. 103(a).

3. Claims 1, 3-5, 43, 45-47, 57, and 59-61 are rejected under 35 U.S.C. 103(a) as being unpatentable over Gleeson et al. (US Patent No: 5,959,989) in view of BECK et al. (US Pub. No: 2001/0014097 A1).

Regarding claim 1, Gleeson et al. teach a method comprising:

receiving a packet, the packet comprising a multicast destination address (see col.13, line 20 wherein a receipt of multicast message is mentioned at MND);

and sending a copy of the packet to a virtual network device sub-unit via a virtual network device link (see col.13, lines 39-48 wherein transmission of multicast message to VLAN designation is mentioned).

Gleeson et al. do not teach specifically the above method wherein the virtual network device link couples two virtual network device sub-units, the two virtual network device sub-units are configured to operate as a single virtual network device, the virtual network device is configured to forward the packet to other layers within a network, and the sending comprises sending at most one copy of the packet from one virtual network device sub-unit to another via the virtual network device link.

However, Beck et al. teach the method wherein the virtual network device link couples two virtual network device sub-units, the two virtual network device sub-units are configured to operate as a single virtual network device (see Fig.7 and also see page 6, para [0064] wherein cluster 24 is shown to include a virtual subnet S3 containing processor nodes A-C and the processor nodes are shown to be coupled by the virtual network device link), the virtual network device is configured to forward the packet to other layers within a network (***see Fig.7 and paragraph [0062], wherein the processor nodes using addresses in different physical subnets sending each other data packets through one or more router nodes is mentioned which is equivalent to the virtual network device configured to forward the packet to other layers within a network*** and also see para [0064]) and the sending comprises sending

at most one copy of the packet from one virtual network device sub-unit to another via the virtual network device link **(**see page 7, paragraphs [0069] & [0070] **wherein sending the packet to one of the processor nodes is mentioned and the receiving processor node transferring the packet to the appropriate/another processor node within the cluster using cluster alias tunneling is mentioned)**.

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the method of Gleeson et al. to include the virtual network device link coupling two virtual network device sub-units and configuring the two virtual network device sub-units to operate as a single virtual network device, the virtual network device being configured to forward the packet to other layers within a network and the sending comprises sending at most one copy of the packet from one virtual network device sub-unit to another via the virtual network device link, disclosed by Beck et al. in order to provide dynamic load distribution in the network and to prevent unnecessary retransmission of data packets in the network and thereby improve the performance of the network for data transmission.

Regarding claim 3, Gleeson et al. further teach the method further comprising: receiving a second packet via the virtual network device link, the second packet comprising a second multicast destination address (see col.15, lines 15-23); and replicating the second packet for each of a plurality of outgoing VLANs (Virtual Local Area Networks) associated with the second multicast destination address (see col.15,

lines 26-28 wherein generation of one or more frames based on MVLAN ID is mentioned).

Regarding claims 4 and 5, Gleeson et al. further teach the method further comprising:

sending at least one copy of the second packet to each line card that includes an interface associated with one of the outgoing VLANs (see Fig.2A and col.15, lines 37-39 wherein sending of the message onto its single port physical interface on trunk 234 by MND 228 is mentioned and trunk 234 which is both incoming/outgoing trunk and is outgoing in this case) and sending at least one copy of the second packet to each line card that includes an interface associated with an incoming VLAN, wherein the second packet is being conveyed in the incoming VLAN (see Fig.2A and col.15, lines 37-39 wherein sending of the message onto its single port physical interface on trunk 234 by MND 228 is mentioned and trunk 234 which is both incoming/outgoing trunk and is incoming in this case).

Regarding claim 43, Gleeson et al. teach a system comprising: means for receiving a packet, the packet comprising a multicast destination address (see col.13, line 20 wherein a receipt of multicast message is mentioned at MND); and means for sending a copy of the packet to a virtual network device sub-unit via a virtual network device link (see col.13, lines 39-48 wherein transmission of multicast message to VLAN designations is mentioned).

Gleeson et al. do not teach specifically the system comprising the virtual network device link couples two virtual network device sub-units, the two virtual network device sub-units are configured to operate as a single virtual network device, the virtual network device is configured to forward the packet to other layers within a network and the means for sending comprises sending at most one copy of the packet from one virtual network device sub-unit to another via the virtual network device link.

However, Beck et al. teach the system comprising the virtual network device link couples two virtual network device sub-units, the two virtual network device sub-units are configured to operate as a single virtual network device (see Fig.7 and also see page 6, para [0064] wherein **cluster** 24 is shown to include a virtual subnet S3 containing processor nodes A-C and the processor nodes are shown to be coupled by the virtual network device link), the virtual network device is configured to forward the packet to other layers within a network (***see Fig.7 and paragraph [0062], wherein the processor nodes using addresses in different physical subnets sending each other data packets through one or more router nodes is mentioned which is equivalent to the virtual network device configured to forward the packet to other layers within a network*** and also see para [0064]) and the means for sending comprises sending at most one copy of the packet from one virtual network device sub-unit to another via the virtual network device link (see page 7, paragraphs [0069] & [0070] **wherein sending the packet to one of the processor nodes is mentioned and the receiving processor node transferring the packet to the**

appropriate/another processor node within the cluster using cluster alias tunneling is mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the system of Gleeson et al. to include the virtual network device link coupling two virtual network device sub-units and configuring the two virtual network device sub-units to operate as a single virtual network device, the virtual network device being configured to forward the packet to other layers within a network and to include the means for sending comprising sending at most one copy of the packet from one virtual network device sub-unit to another via the virtual network device link, disclosed by Beck et al. in order to provide dynamic load distribution in the network and to prevent unnecessary retransmission of data packets in the network and thereby improve the performance of the network for data transmission.

Regarding claim 45, Gleeson et al. further teach the system further comprising: means for receiving a second packet via the virtual network device link, the second packet comprising a second multicast destination address (see col.15, lines 15-23); and means for replicating the second packet for each of a plurality of outgoing VLANs (Virtual Local Area Networks) associated with the second multicast destination address (see col.15, lines 26-28 wherein generation of one or more frames based on MVLAN ID is mentioned).

Regarding claims 46 and 47, Gleeson et al. further teach the system further comprising: means for sending at least one copy of the second packet to each line card that includes an interface associated with one of the outgoing VLANs (see Fig.2A and col.15, lines 37-39 wherein sending of the message onto its single port physical interface on trunk 234 by MND 228 is mentioned and trunk 234 which is both incoming/outgoing trunk and is outgoing in this case) and means for sending at least one copy of the second packet to each line card that includes an interface associated with an incoming VLAN, wherein the second packet is being conveyed in the incoming VLAN (see Fig.2A and col.15, lines 37-39 wherein sending of the message onto its single port physical interface on trunk 234 by MND 228 is mentioned and trunk 234 which is both incoming/outgoing trunk and is incoming in this case).

Regarding claim 57, Gleeson et al. teach a computer readable medium storing a program, the program comprising program instructions executable to (see col.9, lines 50-55): detect reception of a packet, the packet comprising a multicast destination address (see col.13, line 20 wherein a receipt of multicast message is mentioned at MND); and send a copy of the packet to a virtual network device sub-unit via a virtual network device link (see col.13, lines 39-48 wherein transmission of multicast message to VLAN designations is mentioned).

Gleeson et al. do not teach specifically execution of the above computer readable medium wherein having the virtual network device link couples two virtual network device sub-units, the two virtual network device sub-units are configured to

operate as a single virtual network device, the virtual network device is configured to forward the packet to other layers within a network and sending comprises sending at most one copy of the packet from one virtual network device sub-unit to another via the virtual network device link.

However, Beck et al. teach the execution of the computer readable medium wherein having the virtual network device link couples two virtual network device sub-units, and the two virtual network device sub-units are configured to operate as a single virtual network device **(**see Fig.7 and also see page 6, para [0064] wherein **cluster** 24 is shown to include a virtual subnet S3 containing processor nodes A-C and the processor nodes are shown to be coupled by the virtual network device link**),** the virtual network device is configured to forward the packet to other layers within a network (***see Fig.7 and paragraph [0062], wherein the processor nodes using addresses in different physical subnets sending each other data packets through one or more router nodes is mentioned which is equivalent to the virtual network device configured to forward the packet to other layers within a network*** and also see para [0064]) and sending comprises sending at most one copy of the packet from one virtual network device sub-unit to another via the virtual network device link (see page 7, paragraphs [0069] & [0070] **wherein sending the packet to one of the processor nodes is mentioned and the receiving processor node transferring the packet to the appropriate/another processor node within the cluster using cluster alias tunneling is mentioned**).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the execution of the computer readable medium of Gleeson et al. to include the virtual network device link coupling two virtual network device sub-units and configuring the two virtual network device sub-units to operate as a single virtual network device, the virtual network device being configured to forward the packet to other layers within a network and sending comprising sending at most one copy of the packet from one virtual network device sub-unit to another via the virtual network device link, disclosed by Beck et al. in order to provide dynamic load distribution in the network and to prevent unnecessary retransmission of data packets in the network and thereby improve the performance of the network for data transmission.

Regarding claim 59, Gleeson et al. further teach the computer readable medium of claim 58, wherein the program instructions are further executable to: detect reception of a second packet via the virtual network device link, the second packet comprising a second multicast destination address (see col.15, lines 15-23); and replicate the second packet for each of a plurality of outgoing VLANs (Virtual Local Area Networks) associated with the second multicast destination address (see col.15, lines 26-28 wherein generation of one or more frames based on MVLAN ID is mentioned).

Regarding claims 60 and 61, Gleeson et al. further teach the computer readable medium of claim 59, wherein the program instructions are further executable to: send at least one copy of the second packet to each line card that includes an interface associated with one of the outgoing VLANs (see Fig.2A and col.15, lines 37-39

wherein sending of the message onto its single port physical interface on trunk 234 by MND 228 is mentioned and trunk 234 which is both incoming/outgoing trunk and is outgoing in this case) and send at least one copy of the second packet to each line card that includes an interface associated with an incoming VLAN, wherein the second packet is being conveyed in the incoming VLAN (see Fig.2A and col.15, lines 37-39 wherein sending of the message onto its single port physical interface on trunk 234 by MND 228 is mentioned and trunk 234 which is both incoming/outgoing trunk and is incoming in this case).

4. Claims 13-17, 50-54, and 64-68 are rejected under 35 U.S.C. 103(a) as being unpatentable over Kalkunte et al. (US Pub. No: 2003/0198231 A1) in view of BECK et al. (US Pub. No: 2001/0014097 A1).

Regarding claim 13, Kalkunte et al. teach a method, comprising: receiving a packet via a virtual network device link, the packet comprising a unicast destination address (see page 3, para [0037], lines 1-3 wherein receipt of unicast packet by fabric ingress is mentioned and also see page 2, para [0032], lines 6-9 wherein support of VLANs for unicast/broadcast by the fabric is mentioned);

and performing an egress lookup for the packet in response to the receiving the packet (see page 3, para [0037], lines 1-6 wherein **validity of egress port and destination module ID in the header and forwarding of packet to the egress port** are mentioned which is equivalent to performing an egress lookup for the packet in response to the receiving the packet),

wherein the performing the egress lookup comprises allocating a non-primary entry corresponding to a source address of the packet in the lookup table (see page 3, para [0037], lines 7-10 wherein more than one path to destination module in the fabric is mentioned which is equivalent to having a non-primary entry corresponding to the unicast destination address and choosing another egress port based on the ingress port is mentioned).

Kalkunte et al. do not teach specifically the method comprising the virtual network device link couples two virtual network device sub-units, and the two virtual network device sub-units are configured to operate as a single virtual network device.

However, BECK et al. teach the method comprising the virtual network device link couples two virtual network device sub-units and the two virtual network device sub-units are configured to operate as a single virtual network device (see Fig.7 and also see page 6, para [0064] wherein cluster 24 is shown to include a virtual subnet S3 containing processor nodes A-C and the processor nodes are shown to be coupled by the virtual network device link).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the method of Kalkunte et al. to include the virtual network device link coupling two virtual network device sub-units and configuring the

two virtual network device sub-units to operate as a single virtual network device, disclosed by Beck et al. in order to provide dynamic load distribution in the network switch and to prevent unnecessary retransmission of data packets in the network and thereby improve the performance of the network switch for data transmission.

Regarding claim 14, Kalkunte et al. further teach the method wherein a header associated with the packet is also received via the virtual network device link, the header comprises a destination identifier (see page 3, para [0037], lines 1-6 wherein the receipt of header associated with the packet and the destination module id information in the header are mentioned).

Regarding claim 15, Kalkunte et al. further teach the method further comprising: sending the packet and the header to another line card if a non-primary entry corresponding to the unicast destination address is found during the egress lookup (see page 3, para [0037], lines 7-10 wherein more than one path to destination module in the fabric is mentioned which is equivalent to having a non-primary entry corresponding to the unicast destination address and choosing another egress port based on the ingress port is mentioned).

Regarding claim 16, Kalkunte et al. further teach the method further comprising: if a primary entry corresponding to the unicast destination address is found during the egress lookup: sending the packet from an interface identified by the primary entry (see page 3, para [0037], lines 10-14 wherein direct connection of destination modules to the fabric is mentioned which is equivalent to having a primary entry corresponding to the

unicast destination address and the selection of the fabric egress port for sending the packet based on the destination module id and independent of fabric ingress port is mentioned).

Regarding claim 17, Kalkunte et al. further teach the method further comprising: sending a notification via the virtual network device link if the destination identifier comprised in the header does not match a destination identifier comprised in the primary entry, wherein the notification identifies the unicast destination address as corresponding to the destination identifier comprised in the primary entry (see page 3, para [0041] wherein a packet with unknown (Domain Lookup Failure) unicast address is mentioned and in this event, the use of VLAN ID to indicate all the ports the packet is supposed to be delivered is mentioned).

Regarding claim 50, Kalkunte et al. teach a system comprising: means for receiving a packet via a virtual network device link, the packet comprising a unicast destination address (see page 3, para [0037], lines 1-3 wherein receipt of unicast packet by fabric ingress is mentioned and also see page 2, para [0032], lines 6-9 wherein support of VLANs for unicast/broadcast by the fabric is mentioned); and means for performing an egress lookup for the packet (see page 3, para[0037], lines 1-6 wherein validity of egress port in the header and forwarding of packet to the egress port are mentioned),

wherein the means for performing the egress lookup comprises means for allocating a non-primary entry corresponding to a source address of the packet in the lookup table (see page 3, para [0037], lines 7-10 wherein more than one path to

destination module in the fabric is mentioned which is equivalent to having a non-primary entry corresponding to the unicast destination address and choosing another egress port based on the ingress port is mentioned).

.

Kalkunte et al. do not teach specifically the system comprising the virtual network device link couples two virtual network device sub-units, and the two virtual network device sub-units are configured to operate as a single virtual network device.

However, Beck et al. teach the system comprising the virtual network device link couples two virtual network device sub-units and the two virtual network device sub-units are configured to operate as a single virtual network device (see Fig.7 and also see page 6, para [0064] wherein cluster 24 is shown to include a virtual subnet S3 containing processor nodes A-C and the processor nodes are shown to be coupled by the virtual network device link).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the system of Kalkunte et al. to include the virtual network device link coupling two virtual network device sub-units and configuring the two virtual network device sub-units to operate as a single virtual network device, disclosed by Beck et al. in order to provide dynamic load distribution in the network

switch and to prevent unnecessary retransmission of data packets in the network and thereby improve the performance of the network switch for data transmission.

Regarding claim 51, Kalkunte et al. further teach the system wherein a header associated with the packet is also received via the virtual network device link, the header comprises a destination identifier obtained by performing an ingress lookup for the packet (see page 3, para [0037], lines 1-6 wherein the receipt of frame/packet with the header by fabric ingress and the destination module id information in the header are mentioned).

Regarding claim 52, Kalkunte et al. further teach the system further comprising: means for sending the packet and the header to another line card if a non-primary entry corresponding to the unicast destination address is found during the egress lookup (see page 3, para [0037], lines 7-10 wherein more than one path to destination module in the fabric is mentioned which is equivalent to having a non-primary entry corresponding to the unicast destination address and choosing another egress port based on the ingress port is mentioned).

Regarding claim 53, Kalkunte et al. further teach the system further comprising: means for sending the packet from an interface identified by a primary entry, if the primary entry corresponding to the unicast destination address is found during the egress lookup (see page 3, para [0037], lines 10-14 wherein direct connection of destination modules to the fabric is mentioned which is equivalent to having a primary

entry corresponding to the unicast destination address and the selection of the fabric egress port for sending the packet based on the destination module id and independent of fabric ingress port is mentioned).

Regarding claim 54, Kalkunte et al. further teach the system further comprising: means for sending a notification via the virtual network device link if the destination identifier comprised in the header does not match a destination identifier comprised in the primary entry, wherein the notification identifies the unicast destination address as corresponding to the destination identifier comprised in the primary entry (see page 3, para [0041] wherein a packet with unknown (Domain Lookup Failure) unicast address is mentioned and in this event, the use of VLAN ID to indicate all the ports the packet is supposed to be delivered is mentioned).

Regarding claim 64, Kalkunte et al. teach a computer readable medium storing a program, the program comprising program instructions executable to (see page 2, para [0032]): detect reception of a packet via a virtual network device link, the packet comprising a unicast destination address (see page 3, para [0037], lines 1-3 wherein receipt of unicast packet by fabric ingress is mentioned and also see page 2, para [0032], lines 6-9 wherein support of VLANs for unicast/broadcast by the fabric is mentioned); and perform an egress lookup for the packet (see page 3, para[0037], lines 1-6 wherein validity of egress port in the header and forwarding of packet to the egress port are mentioned),

wherein performing the egress lookup comprises allocating a non-primary entry corresponding to a source address of the packet in the lookup table (see page 3,

para [0037], lines 7-10 wherein more than one path to destination module in the fabric is mentioned which is equivalent to having a non-primary entry corresponding to the unicast destination address and choosing another egress port based on the ingress port is mentioned).

Kalkunte et al. do not teach specifically execution of the above computer readable medium having the virtual network device link couples two virtual network device sub-units, and the two virtual network device sub-units are configured to operate as a single virtual network device.

However, Beck et al. teaches the execution of the computer readable medium having the virtual network device link couples two virtual network device sub-units, and the two virtual network device sub-units are configured to operate as a single virtual network device (see Fig.7 and also see page 6, para [0064] wherein cluster 24 is shown to include a virtual subnet S3 containing processor nodes A-C and the processor nodes are shown to be coupled by the virtual network device link).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the execution of the computer readable medium of Kalkunte et al. to include the virtual network device link coupling two virtual network device sub-units and configuring the two virtual network device sub-units to operate as

a single virtual network device, disclosed by Beck et al. in order to provide dynamic load distribution in the network switch and to prevent unnecessary retransmission of data packets in the network and thereby improve the performance of the network switch for data transmission.

Regarding claim 65, Kalkunte et al. further teach the computer readable medium wherein a header associated with the packet is also received via the virtual network device link, the header comprises a destination identifier (see page 3, para [0037], lines 1-6 wherein the receipt of header associated with the packet and the destination module id information in the header are mentioned).

Regarding claim 66, Kalkunte et al. further teach the computer readable medium wherein the program instructions are further executable to send the packet and the header to another line card if a non-primary entry corresponding to the unicast destination address is found during the egress lookup (see page 3, para [0037], lines 7-10 wherein more than one path to destination module in the fabric is mentioned which is equivalent to having a non-primary entry corresponding to the unicast destination address and choosing another egress port based on the ingress port is mentioned).

Regarding claim 67, Kalkunte et al. further teach the computer readable medium wherein the program instructions are further executable to send the packet from an interface identified by a primary entry, if the primary entry corresponding to the unicast destination address is found during the egress lookup (see page 3, para [0037], lines 10-14 wherein direct connection of destination modules to the fabric is mentioned

which is equivalent to having a primary entry corresponding to the unicast destination address and the selection of the fabric egress port for sending the packet based on the destination module id and independent of fabric ingress port is mentioned).

Regarding claim 68, Kalkunte et al. further teach the computer readable medium wherein the program instructions are further executable to send a notification via the virtual network device link if the destination identifier comprised in the header does not match a destination identifier comprised in the primary entry, wherein the notification identifies the unicast destination address as corresponding to the destination identifier comprised in the primary entry (see page 3, para [0041] wherein a packet with unknown (Domain Lookup Failure) unicast address is mentioned and in this event, the use of VLAN ID to indicate all the ports the packet is supposed to be delivered is mentioned).

5. Claims 6-7, 48, and 62 are rejected under 35 U.S.C. 103(a) as being unpatentable over Gleeson et al. (US Patent No: 5,959,989) in view of BECK et al. (US Pub. No: 2001/0014097 A1) and further in view of Ellis et al. (US Pub. No: 2002/0126671 A1).

Regarding claim 6, Gleeson et al. and Beck et al. do not teach specifically the method further comprising: sending at most one copy of the second packet to each line card that includes an interface associated with one of the outgoing VLANs.

However, Ellis et al. teach the method comprising sending at most one copy of the second packet to each line card that includes an interface associated with one of the outgoing VLANs (see page 6, para [0076], lines 1-4 wherein sending each replicated packet according to destination into a VOQ which exists at each ingress/egress port of a fabric card is mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the method of Gleeson et al. and Beck et al. to include sending at most one copy of the second packet to each line card that includes an interface associated with one of the outgoing VLANs disclosed by Ellis et al. for optimum load balancing and faster transmission of multicast packet in the network.

Regarding claim 7, Gleeson et al. and Beck et al. do not teach specifically the method further comprising: not sending any copy of the second packet via an uplink interface coupled to a virtual network device bundle.

However, Ellis et al. teach the method comprising not sending any copy of the second packet via an uplink interface coupled to a virtual network device bundle (see Fig.8 and page 10, para [0109] wherein sending of multicast packet by fabric 811 to ports on LC 806 which is equivalent to sending to downlink interface ports and not sending multicast packet by fabric 811 to ports of LCs 802-804 which is equivalent to uplink interface coupled to a virtual network device bundle are mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the method of Gleeson et al. and Beck et al. to include not sending any copy of the second packet via an uplink interface coupled to a virtual network device bundle disclosed by Ellis et al. for optimum load balancing and faster transmission of multicast packet in the network.

Regarding claim 48, Gleeson et al. and Beck et al. do not teach specifically the system further comprising: means for sending at most one copy of the second packet to each line card that includes an interface associated with one of the outgoing VLANs.

However, Ellis et al. teach the system comprising means for sending at most one copy of the second packet to each line card that includes an interface associated with one of the outgoing VLANs (see page 6, para [0076], lines 1-4 wherein sending each replicated packet according to destination into a VOQ which exists at each ingress/egress port of a fabric card is mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the system of Gleeson et al. and Beck et al. to include means for sending at most one copy of the second packet to each line card that includes an interface associated with one of the outgoing VLANs disclosed by Ellis et al. for optimum load balancing and faster transmission of multicast packet in the network.

Regarding claim 62, Gleeson et al. and Beck et al. do not teach specifically the computer readable medium, wherein the program instructions are further executable to:

send at most one copy of the second packet to each line card that includes an interface associated with one of the outgoing VLANs.

However, Ellis et al. teach the computer readable medium, wherein the program instructions are further executable to send at most one copy of the second packet to each line card that includes an interface associated with one of the outgoing VLANs (see page 6, para [0076], lines 1-4 wherein sending each replicated packet according to destination into a VOQ which exists at each ingress/egress port of a fabric card is mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the computer readable medium of Gleeson et al. and Beck et al., wherein the program instructions are further executable to send at most one copy of the second packet to each line card that includes an interface associated with one of the outgoing VLANs disclosed by Ellis et al. for optimum load balancing and faster transmission of multicast packet in the network.

6. Claims 8-12, 49, and 63 are rejected under 35 U.S.C. 103(a) as being unpatentable over Gleeson et al. (US Patent No: 5,959,989) in view of BECK et al. (US Pub. No: 2001/0014097 A1) and further in view of Kalkunte et al. (US Pub No: 2003/0198231 A1).

Regarding claim 8, Gleeson et al. and Beck et al. do not teach specifically the method further comprising: receiving a third packet via the virtual network device link,

the third packet comprising a unicast destination address; and performing an egress lookup for the third packet in response to the receiving the third packet.

However, Kalkunte et al. teach the method comprising: receiving a third packet via the virtual network device link, the third packet comprising a unicast destination address (see page 3, para [0037], lines 1-3 wherein receipt of unicast packet by fabric ingress is mentioned and also see page 2, para [0032], lines 6-9 wherein support of VLANs for unicast/broadcast by the fabric is mentioned); and performing an egress lookup for the third packet in response to the receiving the third packet (see page 3, para[0037], lines 1-6 wherein validity of egress port in the header and forwarding of packet to the egress port are mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the method of Gleeson et al. and Beck et al. to include receiving a third packet via the virtual network device link, the third packet comprising a unicast destination address and performing an egress lookup for the third packet in response to the receiving the third packet disclosed by Kalkunte et al. to support unicast transmission of packet along with multicast transmission in the network and to improve the performance of the networking system for overall data transmission.

Regarding claim 9, Kalkunte et al. further teach the method wherein a header associated with the third packet is also received via the virtual network device link, the header comprises a destination identifier (see page 3, para [0037], lines 1-6 wherein the

receipt of header associated with the packet and the destination module id information in the header are mentioned).

Regarding claim 10, Kalkunte et al. further teach the method further comprising: sending the third packet and the header to another line card if a non-primary entry corresponding to the unicast destination address is found during the egress lookup (see page 3, para [0037], lines 7-10 wherein more than one path to destination module in the fabric is mentioned which is equivalent to having a non-primary entry corresponding to the unicast destination address and choosing another egress port based on the ingress port is mentioned).

Regarding claim 11, Kalkunte et al. further teach the method further comprising: if a primary entry corresponding to the unicast destination address is found during the egress lookup: sending the third packet from an interface identified by the primary entry (see page 3, para [0037], lines 10-14 wherein direct connection of destination modules to the fabric is mentioned which is equivalent to having a primary entry corresponding to the unicast destination address and the selection of the fabric egress port for sending the packet based on the destination module id and independent of fabric ingress port is mentioned).

Regarding claim 12, Kalkunte et al. further teach the method further comprising: sending a notification via the virtual network device link if the destination identifier comprised in the header does not match a destination identifier comprised in the primary entry, wherein the notification identifies the unicast destination address as corresponding to the destination identifier comprised in the primary entry (see page 3,

para [0041] wherein a packet with unknown (Domain Lookup Failure) unicast address is mentioned and in this event, the use of VLAN ID to indicate all the ports the packet is supposed to be delivered is mentioned).

Regarding claim 49, Gleeson et al. and Beck et al. do not teach specifically the system further comprising: means for receiving a third packet via the virtual network device link, the third packet comprising a unicast destination address; and means for performing an egress lookup for the third packet in response to receiving the third packet.

However, Kalkunte et al. teach the system comprising: means for receiving a third packet via the virtual network device link, the third packet comprising a unicast destination address (see page 3, para [0037], lines 1-3 wherein receipt of unicast packet by fabric ingress is mentioned and also see page 2, para [0032], lines 6-9 wherein support of VLANs for unicast/broadcast by the fabric is mentioned); and means for performing an egress lookup for the third packet in response to receiving the third packet (see page 3, para[0037], lines 1-6 wherein validity of egress port in the header and forwarding of packet to the egress port are mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the system of Gleeson et al. and Beck et al. to include means for receiving a third packet via the virtual network device link, the third packet comprising a unicast destination address and means for performing an egress lookup

for the third packet in response to receiving the third packet disclosed by Kalkunte et al. to support unicast transmission of packet along with multicast transmission in the network and for optimum load balancing of traffic in the network.

Regarding claim 63, Gleeson et al. and Beck et al. do not teach specifically the computer readable medium, wherein the program instructions are further executable to: detect reception of a third packet via the virtual network device link, the third packet comprising a unicast destination address; and perform an egress lookup for the third packet in response to the receiving the third packet.

However, Kalkunte et al. teach the computer readable medium, wherein the program instructions are further executable to detect reception of a third packet via the virtual network device link, the third packet comprising a unicast destination address (see page 3, para [0037], lines 1-3 wherein receipt of unicast packet by fabric ingress is mentioned and also see page 2, para [0032], lines 6-9 wherein support of VLANs for unicast/broadcast by the fabric is mentioned) and perform an egress lookup for the third packet in response to the receiving the third packet (see page 3, para[0037], lines 1-6 wherein validity of egress port in the header and forwarding of packet to the egress port are mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the computer readable medium of Gleeson et al. and Beck et al. to modify the executable program instructions to detect reception of a third

packet via the virtual network device link, the third packet comprising a unicast destination address and perform an egress lookup for the third packet in response to the receiving the third packet disclosed by Kalkunte et al. to support unicast transmission of packet along with multicast transmission in the network and for optimum load balancing of traffic in the network.

7. Claims 34, 36-37, and 39 are rejected under 35 U.S.C. 103(a) as being unpatentable over Gleeson et al. (US Patent No: 5,959,989) in view of BECK et al. (US Pub. No: 2001/0014097 A1), further in view of Gallo et al. (US Patent No: 6,760,776 B1).

Regarding claim 34, Gleeson et al. teach a system comprising: an interface to a virtual network device link, wherein the interface is configured to receive a packet (see col.13, line 20 wherein a receipt of multicast message is mentioned at MND);

and a distributed forwarding module coupled to the interface, wherein the distributed forwarding module is configured to forward the packet (see col.13, lines 39-48 wherein distribution of multicast message by multicast controller of MND to VLAN designations is mentioned).

Gleeson et al. do not teach specifically the system comprising the virtual network device link couples two virtual network device sub- units, and the two virtual network device sub-units are configured to operate as a single virtual network device.

However, BECK et al. teach the system comprising the virtual network device link couples two virtual network device sub- units, and the two virtual network device sub-units are configured to operate as a single virtual network device (see Fig.7 and also see page 6, para [0064] wherein **cluster** 24 is shown to include a virtual subnet S3 containing processor nodes A-C and the processor nodes are shown to be coupled by the virtual network device link).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the system of Gleeson et al. to include the virtual network device link coupling two virtual network device sub-units and configuring the two virtual network device sub-units to operate as a single virtual network device, disclosed BECK et al. by in order to provide dynamic load distribution in the network and to prevent unnecessary retransmission of data packets in the network and thereby improve the performance of the network for data transmission.

Gleeson et al. and BECK et al. do not yet together teach specifically the system wherein the distributed forwarding module is configured to perform an ingress lookup for the packet if the packet includes a multicast destination address, and the distributed forwarding module is configured to perform an egress lookup for the packet if the packet includes a unicast destination address.

However, Gallo et al. teach the system wherein the distributed forwarding module is configured to perform an ingress lookup for the packet if the packet includes a multicast destination address (see Fig.4, and col.3, lines 47-59) and the distributed forwarding module is configured to perform an egress lookup for the packet if the packet includes a unicast destination address (see Fig.4, and col.4, lines 24-26 and lines 31-34).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the system of Gleeson et al. and BECK et al. to include the configuration of the distributed forwarding module to perform an ingress lookup for the packet if the packet includes a multicast destination address and to include the configuration of the distributed forwarding module to perform an egress lookup for the packet if the packet includes a unicast destination address, disclosed by Gallo et al. to support both multicast transmission and unicast transmission of the packet in the network and to further improve the overall performance of the system for both multicast and unicast data transmission.

Regarding claim 36, Gleeson et al. further teach the system wherein the packet includes a multicast destination address (see col.15, lines 15-23), and the distributed forwarding module is configured to replicate the packet for each of a plurality of outgoing VLANs associated with the multicast destination address (see col.15, lines 26-28 wherein generation of one or more frames based on MVLAN ID by multicast controller at MND is mentioned).

Regarding claim 37, Gleeson et al. further teach the system further comprising: one or more line cards, wherein the distributed forwarding module is configured to send at least one copy of the packet to each of the one or more line cards that includes an interface associated with one of the outgoing VLANs (see Fig.2A and col.15, lines 37-39 wherein sending of the message onto its single port physical interface on trunk 234 by MND 228 is mentioned and trunk 234 which is both incoming/outgoing trunk and is outgoing in this case).

Regarding claim 39, Gleeson et al. further teach the system further comprising: a second interface configured to receive a second packet, wherein the second packet comprises a second multicast address (see col.13, line 20 wherein a receipt of multicast message is mentioned at MND), and the distributed forwarding module is configured to send at most one copy of the second packet via the virtual network device link (see col.13, lines 46-53 and see col.15, lines 6-12 wherein sending of at most one copy of the message to VLAN segment of the network is mentioned).

8. Claims 40-42 are rejected under 35 U.S.C. 103(a) as being unpatentable over Gleeson et al. (US Patent No: 5,959,989) in view of BECK et al. (US Pub. No: 2001/0014097 A1), further in view of Gallo et al. (US Patent No: 6,760,776 B1) and further in view of Kalkunte et al. (US Pub No: 2003/0198231 A1).

Regarding claim 40, Gleeson et al., Beck et al. and Gallo et al. do not teach specifically the system, wherein a header associated with the packet is also received via

the virtual network device link, the header comprises a destination identifier, and the packet comprises the unicast destination address, and the distributed forwarding module is configured to send the packet and the header to another line card if a non-primary entry corresponding to the unicast destination address is found during the egress lookup.

However, Kalkunte et al. teach the system, wherein a header associated with the packet is also received via the virtual network device link, the header comprises a destination identifier (see page 3, para[0037], lines 1-6 wherein the receipt of header associated with the packet and the destination module id information in the header are mentioned), and the packet comprises the unicast destination address (see page 3, para [0037], lines 1-3 wherein receipt of unicast packet by fabric ingress is mentioned and also see page 2, para [0032], lines 6-9 wherein support of VLANs for unicast/broadcast by the fabric is mentioned), and the distributed forwarding module is configured to send the packet and the header to another line card if a non-primary entry corresponding to the unicast destination address is found during the egress lookup (see page 3, para [0037], lines 7-10 wherein more than one path to destination module in the fabric is mentioned which is equivalent to having a non-primary entry corresponding to the unicast destination address and choosing another egress port based on the ingress port is mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the system of Gleeson et al., Beck et al. and Gallo et al. to include the receipt of a packet along with the header via the virtual network device link, the header comprises a destination identifier, and the packet comprises the unicast destination address and the distributed forwarding module is configured to send the packet and the header to another line card if a non-primary entry corresponding to the unicast destination address is found during the egress lookup disclosed by Kalkunte et al. to support unicast transmission of packet along with multicast transmission and for optimum load balancing of traffic in the network.

Regarding claim 41, Kalkunte et al. further teach the system further comprising: a second interface, wherein the distributed forwarding module is configured to send the packet from the second interface if a primary entry corresponding to the unicast destination address is found during the egress lookup and if the primary entry identifies the second interface (see page 3, para [0037], lines 10-14 wherein direct connection of destination modules to the fabric is mentioned which is equivalent to having a primary entry corresponding to the unicast destination address and the selection of the fabric egress port for sending the packet based on the destination module id and independent of fabric ingress port is mentioned).

Regarding claim 42, Kalkunte et al. further teach the system wherein the distributed forwarding module is configured to send a notification via the virtual network device link if a destination identifier comprised in the header does not match a destination identifier comprised in the primary entry, and the notification identifies the

unicast destination address as corresponding to the destination identifier comprised in the primary entry (see page 3, para [0041] wherein a packet with unknown (Domain Lookup Failure) unicast address is mentioned and in this event, the use of VLAN ID to indicate all the ports the packet is supposed to be delivered is mentioned).

9. Claim 38 is rejected under 35 U.S.C. 103(a) as being unpatentable over Gleeson et al. (US Patent No: 5,959,989) in view of BECK et al. (US Pub. No: 2001/0014097 A1), further in view of Gallo et al. (US Patent No: 6,760,776 B1) and further in view of Ellis et al. (US Pub No: 2002/0126671 A1).

Regarding claim 38, Gleeson et al., Beck et al. and Gallo et al. do not teach specifically the system further comprising: one or more line cards, wherein the distributed forwarding module is configured to send at most one copy of the packet to each line card that includes an interface associated with one of the outgoing VLANs.

However, Ellis et al. teach the system comprising one or more line cards, wherein the distributed forwarding module is configured to send at most one copy of the packet to each line card that includes an interface associated with one of the outgoing VLANs (see page 6, para [0076], lines 1-4 wherein sending each replicated packet according to destination into a VOQ which exists at each ingress/egress port of a fabric card is mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the system of Gleeson et al., Beck et al. and Gallo et al. to include the configuration of the distributed forwarding module to send at most one copy of the packet to each line card that includes an interface associated with one of the outgoing VLANs disclosed by Ellis et al. for optimum load balancing and faster transmission of multicast packet in the network.

10. Claims 18, 55, and 69 are rejected under 35 U.S.C. 103(a) as being unpatentable over Kalkunte et al. (US Pub No: 2003/0198231 A1) in view of BECK et al. (US Pub. No: 2001/0014097 A1) and further in view of Ellis et al. (US Pub No: 2002/0126671 A1).

Regarding claims 18, 55 and 69, Kalkunte et al. and Beck et al. do not teach specifically the method/system/computer readable medium wherein the packet is only sent from the interface if the interface is not comprised in an uplink interface bundle.

However, Ellis et al. teach the method/system/computer readable medium wherein the packet is only sent from the interface if the interface is not comprised in an uplink interface bundle (see Fig.8 and page 10, para [0109] wherein sending of multicast packet by fabric 811 to ports on LC 806 which is equivalent to sending to downlink interface ports and not sending multicast packet by fabric 811 to ports of LCs 802-804 which is equivalent to uplink interface bundle are mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the method/system/computer readable medium of Kalkunte et al. and beck et al. to include sending the packet from the interface if the interface is not comprised in an uplink interface bundle disclosed by Ellis et al. for optimum load balancing and faster transmission of multicast packet in the network.

11. Claims 19-22, 56, and 70 are rejected under 35 U.S.C. 103(a) as being unpatentable over Kalkunte et al. (US Pub No: 2003/0198231 A1) in view of BECK et al. (US Pub. No: 2001/0014097 A1) and further in view of Gleeson et al. (US Patent No: 5,959,989).

Regarding claim 19, Kalkunte et al. and Beck et al. do not teach specifically the method further comprising: receiving a second packet, the second packet comprising a multicast destination address; and sending at most one copy of the second packet to one of the two virtual network device sub-units via the virtual network device link.

However, Gleeson et al. teach the method further comprising receiving a second packet, the second packet comprising a multicast destination address (see col.13, line 20 wherein a receipt of multicast message is mentioned at MND) and sending at most one copy of the second packet to one of the two virtual network device sub-units via the virtual network device link (see col.13, lines 39-48 wherein transmission of multicast message to VLAN designation is mentioned and also see col.15, lines 6-12 wherein sending of at most one copy of the message to VLAN segment of the network is mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the method of Kalkunte et al. and Beck et al. to include the receipt of a second packet comprising a multicast destination address and sending at most one copy of the second packet to one of the two virtual network device sub-units via the virtual network device link, disclosed by Gleeson et al. for proper load balancing and optimum multicast transmission of packet in the network.

Regarding claim 20, Gleeson et al. further teach the method further comprising: receiving a third packet via the virtual network device link, the third packet comprising a second multicast destination address (see col.15, lines 15-23) and replicating the third packet for each of a plurality of outgoing VLANs (Virtual Local Area Networks) associated with the second multicast destination address (see col.15, lines 26-28 wherein generation of one or more frames based on MVLAN ID is mentioned).

Regarding claim 21, Gleeson et al. further teach the method further comprising: sending at least one copy of the third packet to each line card that includes an interface associated with one of the outgoing VLANs (see Fig.2A and col.15, lines 37-39 wherein sending of the message onto its single port physical interface on trunk 234 by MND 228 is mentioned and trunk 234 which is both incoming/outgoing trunk and is outgoing in this case).

Regarding claim 22, Gleeson et al. further teach the method further comprising: sending at least one copy of the third packet to each line card that includes an interface associated with an incoming VLAN, wherein the third packet is being conveyed in the

incoming VLAN (see Fig.2A and col.15, lines 37-39 wherein sending of the message onto its single port physical interface on trunk 234 by MND 228 is mentioned and trunk 234 which is both incoming/outgoing trunk and is incoming in this case).

Regarding claim 56, Kalkunte et al. and Beck et al. do not teach specifically the system further comprising: means for receiving a second packet, the second packet comprising a multicast destination address; and means for sending at most one copy of the second packet to one of the two virtual network device sub-units via the virtual network device link.

However, Gleeson et al. teach the system further comprising means for receiving a second packet, the second packet comprising a multicast destination address (see col.13, line 20 wherein a receipt of multicast message is mentioned at MND) and means for sending at most one copy of the second packet to one of the two virtual network device sub-units via the virtual network device link (see col.13, lines 39-48 wherein transmission of multicast message to VLAN designation is mentioned and also see col.15, lines 6-12 wherein sending of at most one copy of the message to VLAN segment of the network is mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the system of Kalkunte et al. and Beck et al. to include the means for receiving a second packet comprising a multicast destination address and means for sending at most one copy of the second packet to one of the two virtual

network device sub-units via the virtual network device link, disclosed by Gleeson et al. for proper load balancing and optimum multicast transmission of packet in the network.

Regarding claim 70, Kalkunte et al. and Beck et al. do not teach specifically the computer readable medium wherein the program instructions are further executable to detect reception of a second packet, the second packet comprising a multicast destination address; and send at most one copy of the second packet to a virtual network device sub-unit via a virtual network device link, the virtual network device sub-unit comprised in a virtual network device.

However, Gleeson et al. teach the computer readable medium wherein the program instructions are further executable to detect reception of a second packet, the second packet comprising a multicast destination address (see col.13, line 20 wherein a receipt of multicast message is mentioned at MND) and send at most one copy of the second packet to a virtual network device sub-unit via a virtual network device link, the virtual network device sub-unit comprised in a virtual network device (see col.13, lines 39-48 wherein transmission of multicast message to VLAN designation is mentioned and also see col.15, lines 6-12 wherein sending of at most one copy of the message to VLAN segment of the network is mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the computer readable medium of Kalkunte et al. and Beck et al. to include detecting reception of a second packet comprising a multicast

destination address and sending at most one copy of the second packet to a virtual network device sub-unit via a virtual network device link, the virtual network device sub-unit comprised in a virtual network device disclosed by Gleeson et al. for proper load balancing and optimum multicast transmission of packet in the network.

12. Claim 23 is rejected under 35 U.S.C. 103(a) as being unpatentable over Kalkunte et al. (US Pub No: 2003/0198231 A1) in view of Beck et al. (US Pub. No: 2001/0014097 A1), further in view of Gleeson et al. (US Patent No: 5,959,989) and further in view of Ellis et al. (US Pub. No: 2002/0126671 A1).

Regarding claim 23, Kalkunte et al., Beck et al. and Gleeson et al. do not teach specifically the method further comprising: sending at most one copy of the third packet to each line card that includes an interface associated with one of the outgoing VLANs.

However, Ellis et al. teach the method further comprising sending at most one copy of the third packet to each line card that includes an interface associated with one of the outgoing VLANs (see page 6, para [0076], lines 1-4 wherein sending each replicated packet according to destination into a VOQ which exists at each ingress/egress port of a fabric card is mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the method of Kalkunte et al., Beck et al. and Gleeson et al. to include sending at most one copy of the second packet to each line card that

includes an interface associated with one of the outgoing VLANs disclosed by Ellis et al. for optimum load balancing and faster transmission of multicast packet in the network.

13. Claims 24, and 30-32 are rejected under 35 U.S.C. 103(a) as being unpatentable over Kalkunte et al. (US Pub No: 2003/0198231 A1) in view of Gallo et al. (US Patent No: 6,760,776 B1).

Regarding claim 24, Kalkunte et al. teach a method comprising: receiving a packet via a virtual network device link (see page 3, para [0037], lines 1-3 wherein receipt of a packet by fabric ingress is mentioned and also see page 2, para [0032], lines 6-9 wherein support of VLANs for unicast/broadcast by the fabric is mentioned); performing an egress lookup for the packet, wherein the egress lookup is performed for the packet if the packet includes a unicast destination address **(see page 3, para[0037], lines 1-6 wherein validity of unicast packet and egress port/destination module ID in the header of unicast packet and forwarding of packet to the egress port are mentioned which is equivalent to performing an egress lookup for the packet that includes a unicast destination address),**

wherein the performing the egress lookup comprises allocating a non-primary entry corresponding to a source address of the packet in the lookup table (see page 3, para [0037], lines 7-10 wherein more than one path to destination module in the fabric is mentioned which is equivalent to having a non-primary entry corresponding to the unicast destination address and choosing another egress port based on the ingress port is mentioned).

Kalkunte et al. do not teach specifically the method includes performing one of an ingress lookup for the received packet, wherein the ingress lookup is performed for the packet if the packet includes a multicast destination address and a primary lookup table entry can be allocated in response to an ingress lookup but not in response to an egress lookup.

However, Gallo et al. teach the method includes performing one of an ingress lookup for the received packet, wherein the ingress lookup is performed for the packet if the packet includes a multicast destination address **(see Fig.4 and col.3, lines 47-59 wherein performing the ingress lookup of packet that contains multicast destination address is mentioned)** and a primary lookup table entry can be allocated in response to an ingress lookup but not in response to an egress lookup **(see Fig.4 and col.3, lines 59-63 wherein allocating routing table that contains a special identifier for layer 3 ingress processor is mentioned which is equivalent to allocating a primary lookup table entry in response to an ingress lookup)**.

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the method of Kalkunte et al. to include performing one of an ingress lookup for the received packet, wherein the ingress lookup is performed for the packet if the packet includes a multicast destination address and allocating a primary lookup table entry in response to an ingress lookup, disclosed by Gallo et al. to

obtain proper load balancing of traffic in the network and thereby to improve the performance of data transmission in the network switch.

Regarding claim 30, Kalkunte et al. further teach the method wherein a header associated with the packet is also received via the virtual network device link, the header comprises a destination identifier (see page 3, para[0037], lines 1-6 wherein the receipt of header associated with the packet and the destination module id information in the header are mentioned), and the packet comprises the unicast destination address(see page 3, para [0037], lines 1-3 wherein receipt of unicast packet by fabric ingress is mentioned), and the method further comprises: sending the packet and the header to another line card if a non-primary entry corresponding to the unicast destination address is found during the egress lookup (see page 3, para [0037], lines 7-10 wherein more than one path to destination module in the fabric is mentioned which is equivalent to having a non-primary entry corresponding to the unicast destination address and choosing another egress port based on the ingress port is mentioned).

Regarding claim 31, Kalkunte et al. further teach the method further comprising: if a primary entry corresponding to the unicast destination address is found during the egress lookup: sending the packet from an interface identified by the primary entry (see page 3, para [0037], lines 10-14 wherein direct connection of destination modules to the fabric is mentioned which is equivalent to having a primary entry corresponding to the unicast destination address and the selection of the fabric egress port for sending the packet based on the destination module id and independent of fabric ingress port is mentioned).

Regarding claim 32, Kalkunte et al. further teach the method further comprising: sending a notification via the virtual network device link if a destination identifier comprised in the header does not match a destination identifier comprised in the primary entry, wherein the notification identifies the unicast destination address as corresponding to the destination identifier comprised in the primary entry (see page 3, para [0041] wherein a packet with unknown (Domain Lookup Failure) unicast address is mentioned and in this event, the use of VLAN ID to indicate all the ports the packet is supposed to be delivered is mentioned).

14. Claims 25, 26, and 28 are rejected under 35 U.S.C. 103(a) as being unpatentable over Kalkunte et al. (US Pub No: 2003/0198231 A1) in view of Gallo et al. (US Patent No: 6,760,776 B1) and further in view of Gleeson et al. (US Patent No: 5,959,989).

Regarding claim 25, Kalkunte et al. and Gallo et al. do not teach specifically the method wherein the packet includes a multicast destination address, and the method further comprises: replicating the packet for each of a plurality of outgoing VLANs associated with the multicast destination address.

However, Gleeson et al. teach the method wherein the packet includes a multicast destination address (see col.15, lines 15-23), and the method further comprises: replicating the packet for each of a plurality of outgoing VLANs associated with the multicast destination address (see col.15, lines 26-28 wherein generation of one or more frames based on MVLAN ID is mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the method of Kalkunte et al. and Gallo et al. to include replicating the packet that has multicast destination address for each of a plurality of outgoing VLANs associated with the multicast destination address disclosed by Gleeson et al. for optimum and faster transmission of multicast transmission of packets in the network.

Regarding claim 26, Gleeson et al. further teach the method further comprising: sending at least one copy of the packet to each line card that includes an interface associated with one of the outgoing VLANs (see Fig.2A and col.15, lines 37-39 wherein sending of the message onto its single port physical interface on trunk 234 by MND 228 is mentioned and trunk 234 which is both incoming/outgoing trunk and is outgoing in this case).

Regarding claim 28, Gleeson et al. further teach the method further comprising: not sending any copy of the packet via the virtual network device link (see col.14, lines 26-31 wherein not sending of the message on port five by device 221 is mentioned).

15. Claims 27 and 29 are rejected under 35 U.S.C. 103(a) as being unpatentable over Kalkunte et al. (US Pub No: 2003/0198231 A1) in view of Gallo et al. (US Patent No: 6,760,776 B1), further in view of Gleeson et al. (US Patent No: 5,959,989) and further in view of Ellis et al. (US Pub. No: 2002/0126671 A1).

Regarding claim 27, Kalkunte et al., Gallo et al. and Gleeson et al. do not teach specifically the method further comprising: sending at most one copy of the packet to each line card that includes an interface associated with one of the outgoing VLANs.

However, Ellis et al. teach the method comprising sending at most one copy of the packet to each line card that includes an interface associated with one of the outgoing VLANs (see page 6, para [0076], lines 1-4 wherein sending each replicated packet according to destination into a VOQ which exists at each ingress/egress port of a fabric card is mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the method of Kalkunte et al., Gallo et al. and Gleeson et al. to include sending at most one copy of the second packet to each line card that includes an interface associated with one of the outgoing VLANs disclosed by Ellis et al. to further improve load balancing of traffic and thereby to improve transmission efficiency of multicast data transmission in the network.

Regarding claim 29, Kalkunte et al., Gallo et al. and Gleeson et al. do not teach specifically the method further comprising: not sending any copy of the packet via an uplink interface comprised in an uplink interface bundle.

However, Ellis et al. teach the method comprising not sending any copy of the packet via an uplink interface comprised in a uplink interface bundle (see Fig.8 and

page 10, para [0109] wherein sending of multicast packet by fabric 811 to ports on LC 806 which is equivalent to sending to downlink interface ports and not sending multicast packet by fabric 811 to ports of LCs 802-804 which is equivalent to uplink interface coupled to a virtual network device bundle are mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the method of Kalkunte et al., Gallo et al. and Gleeson et al. to include not sending any copy of the second packet via an uplink interface coupled to a virtual network device bundle disclosed by Ellis et al. to further improve load balancing of traffic and thereby to improve transmission efficiency of multicast data transmission in the network.

16. Claim 33 is rejected under 35 U.S.C. 103(a) as being unpatentable over Kalkunte et al. (US Pub No: 2003/0198231 A1) in view of Gallo et al. (US Patent No: 6,760,776 B1) and further in view of Ellis et al. (US Pub. No: 2002/0126671 A1).

Regarding claim 33, Kalkunte et al. and Gallo et al. do not teach specifically the method wherein the packet is only sent from the interface if the interface is not comprised in an uplink interface bundle.

However, Ellis et al. teach the method wherein the packet is only sent from the interface if the interface is not comprised in a uplink interface bundle (see Fig.8 and page 10, para [0109] wherein sending of multicast packet by fabric 811 to ports on LC 806 which is equivalent to sending to downlink interface ports and not sending multicast

packet by fabric 811 to ports of LCs 802-804 which is equivalent to uplink interface bundle are mentioned).

Therefore, it would have been obvious to one of ordinary skills in the art at the time of the invention to modify the method of Kalkunte et al. and Gallo et al. to include sending the packet via an interface only if the interface is not comprised in a uplink interface bundle disclosed by Ellis et al. for optimum load balancing of traffic and faster transmission of multicast packet in the network.

Response to Arguments

17. Applicant's arguments filed on 06/17/2009 have been fully considered but they are not persuasive. Applicant's amendment of claims necessitated new citations of the references as mentioned above under Claim Rejections.

18. In pages 20-21, regarding claims 1, 43 and 57, Applicant mentions that the cited sections of Beck fail to show, teach, or even suggest that Beck's processing nodes are configured to operate as a single virtual network device and are configured to forward a packet to other layers in a network and there is no teaching or suggestion in Beck that the processor nodes are configured to forward a packet to other layers in a network and therefore, the cited sections of Beck fail to show, teach, or even suggest two virtual network device subunits being configured to forward a packet to other layers in a network.

However, the Examiner respectfully disagrees to the above statements of the Applicant as Beck clearly teaches that the processing nodes are configured to operate as a single virtual network device (see Fig.7 and also see page 6, para [0064] wherein cluster 24 is shown to include a virtual subnet S3 containing processor nodes A-C and the processor nodes are shown to be coupled by the virtual network device link and these processor nodes operate as a single virtual network device) and are configured to forward a packet to other layers in a network (**see Fig.7 and paragraph [0062], wherein the processor nodes using addresses in different physical subnets sending each other data packets *through one or more router nodes* is mentioned which is equivalent to the virtual network device configured to forward the packet to other layers within a network** as the data packets from processor nodes of one physical subnet send packets to the processor nodes of different physical subnet *via one or more network routers* in the network and thus the virtual network device being configured to forward the packet to other layers within a network as these packets pass through *one or more network routers* in the network) and thus Gleeson et al. in combination with Beck teach all the limitations of the claims 1, 43 and 57 as already mentioned above under Claim Rejections.

19. In pages 22-23 of Applicant's Remarks, regarding claims 13, 50 and 64, Applicant mentions that the cited sections of Kalkunte fail to show, teach, or even suggest performing an egress lookup that comprises allocating a non-primary entry corresponding to a source address of a packet in the lookup table and further mentions

that Kalkunte fails to disclose an allocating operation for allocating a non-primary entry in a lookup table, and certainly fails to disclose that such allocation of a non-primary entry is performed as part of an egress lookup.

However, the Examiner respectfully disagrees to the above statements of the Applicant as Kalkunte clearly teaches performing an egress lookup for the packet in response to the receiving the packet (see page 3, para [0037], lines 1-6 wherein ***validity of egress port and destination module ID in the header and forwarding of packet to the egress port*** are mentioned which is equivalent to performing an egress lookup for the packet in response to the receiving the packet and also see paragraph [0038] wherein egress lookup/routing table for unicast packet is mentioned) and performing an egress lookup that comprises allocating a non-primary entry corresponding to a source address of a packet in the lookup table (see page 3, para [0037], lines 7-10 *wherein more than one path to destination module in the fabric is mentioned which is equivalent to having a non-primary entry corresponding to the unicast destination address* and choosing another egress port based on the ingress port/source address of a packet is also mentioned).

20. In pages 23-25 of Applicant's Remarks, regarding claim 24, Applicant mentions that the cited sections of *Kalkunte* fail to show, teach, or even suggest performing an egress lookup that comprises allocating a non-primary entry corresponding to a source address of a packet in the lookup table and further mentions that the cited sections of *Gallo* fail to disclose allocating a primary lookup entry in response to an ingress lookup

and the cited sections of Gallo fail to disclose that a primary lookup table entry is not allocated in response to an egress lookup.

However, the Examiner respectfully disagrees to the above statements of the Applicant as *Kalkunte* teaches performing an egress lookup that comprises allocating a non-primary entry corresponding to a source address of a packet in the lookup table (see page 3, para [0037], lines 7-10 *wherein more than one path to destination module in the fabric is mentioned which is equivalent to having a non-primary entry corresponding to the unicast destination address* and choosing another egress port based on the ingress port/source address of a packet is also mentioned) as already explained above in section 19 and *Gallo* teaches allocating a primary lookup entry in response to an ingress lookup but not in response to an egress lookup (see Fig.4 and col.3, lines 47-59 **wherein performing the ingress lookup of packet that contains multicast destination address is mentioned** and see **col.3, lines 59-63 wherein allocating routing table that contains a special identifier for layer 3 ingress processor is mentioned which is equivalent to allocating a primary lookup table entry in response to an ingress lookup and this primary lookup table entry is not allocated in response to an egress look up as this processing involves layer 3 ingress processor and see col.3, line 64 to col. 4, line 2 wherein, for the packets that are not multicast, the packets are routed to the output media control filter 55, which may be located on the egress NP, is mentioned).**

21. The rejection of all other claims is already explained under Claim Rejections above.

Conclusion

22. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

23. Any response to this office action should be faxed to (571) 273-8300 or mailed To:

Commissioner for Patents,
P.O. Box 1450
Alexandria, VA 22313-1450
Hand-delivered responses should be brought to
Customer Service Window
Randolph Building
401 Dulany Street

Alexandria, VA 22314.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to SRINIVASA R. REDDIVALAM whose telephone number is (571)270-3524. The examiner can normally be reached on Mon-Fri 9:30 AM - 6:30 PM.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Chirag Shah can be reached on 571-272-3144. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

Srini Reddivalam
11/05/2009

/Chirag G Shah/
Supervisory Patent Examiner, Art Unit 2477